

METHOD AND PACKET-LEVEL DEVICE FOR TRAFFIC REGULATION IN A DATA NETWORK

BACKGROUND OF THE INVENTION

1. Field of the Invention

5 The present invention relates to bucket mechanisms for traffic regulation in a data network.

2. Background Art

 A packet communication network accommodates data contained inside packets. Packet traffic, if not regulated, may cause network congestion. During network congestion, packets may be dropped or lost. Traffic regulation techniques attempt to manage network traffic to avoid congestion. Bucket mechanisms are one technique used to regulate traffic at the packet level. Bucket mechanisms include one or more token buckets, leaky buckets, and any other bucket arrangement used for policing and/or shaping data by using tokens or the equivalent thereof.

 An existing traffic regulator using a bucket mechanism includes a token bucket. The token bucket receives new tokens at a fixed rate. The maximum number of tokens that the bucket can hold is the bucket depth. If the bucket is already full when a new token is received, the new token is ignored. When a packet arrives at the regulator, a number of tokens equal to the bit size (or byte size in some applications) of the packet must be removed from the bucket for the packet to be sent normally. In the event that there are not enough tokens present in the bucket, the packet is either dropped or remarked with a lower priority and then sent. The rate that new tokens are received at the token bucket defines the maximum sustained normal data rate, while the bucket depth defines the maximum data burst size. The bucket mechanism may optionally include multiple policy support. That is, different classifications of traffic may be regulated by different token buckets

such that the token bucket for one class of traffic may be defined differently than the token bucket for another class of traffic.

An existing packet communication network that is used extensively is the Internet. The Internet is a collection of interconnected networks. An Internet service provider (ISP) provides its subscribers access to the Internet. The transmission control protocol/Internet protocol (TCP/IP) family of protocols is used for communication over the Internet. A typical application for a traffic regulator using a bucket mechanism is to locate the traffic regulator at the subscriber's connection to the ISP to regulate downstream data flow (from the ISP to the subscriber). The ISP configures the traffic regulator by defining the bucket mechanism parameters. The parameters remain fixed and are configured according to the subscriber agreement to provide the agreed upon maximum sustained normal data rate in the downstream direction. A suitable technique is used to regulate upstream data transfer.

The ISP typically bills the subscriber for Internet access in a flat rate arrangement that allows unlimited access for a fixed monthly or yearly fee. This arrangement simplifies billing but results in a situation where a small portion of the subscriber base is responsible for a majority of the ISP network costs. That is, a small portion of the subscriber base demands a majority of the network resources.

Accordingly, this flat rate billing arrangement for Internet access typically used by ISP's results in a somewhat inequitable distribution of service costs to the subscribers. These inequities could be addressed by modifying the billing approach. For example, a tiered billing structure or strict usage-based billing could be used to better match the amount billed to a subscriber to the actual subscriber usage. However, in certain situations it may be desirable to stay with flat rate billing. Further, there may be certain situations where modifying the billing approach does not fully address the problem of inequitable service cost distribution.

Background information may be found in U.S. Patent Nos. 5,289,462 and 6,147,970, and in Cisco IOS Quality of Service Solutions Configuration Guide.

Policing and Shaping Overview (Cisco Systems, Inc., May 2003), pages QC-203 - QC-218.

5 For the foregoing reasons, there is a need for an improved traffic regulation method that provides a more equitable distribution of network resources to subscribers.

SUMMARY OF THE INVENTION

It is an object of the present invention to regulate traffic in a packet communication network in a way that dynamically adjusts a subscriber associated token bucket based on the resource demand placed on the network by the subscriber.

10 In carrying out the present invention, a method of traffic regulation in a packet communication network is provided. The network includes a traffic regulator for regulating traffic at the packet level. The traffic regulator includes a bucket mechanism. The bucket mechanism includes a token bucket associated with a subscriber. The token bucket is configured to receive new tokens at a fill rate and
15 is configured with a bucket depth. The method comprises handling packets that arrive at the regulator in accordance with the token bucket configuration. The method further comprises measuring a demand placed on the packet communication network by the subscriber. The method further comprises dynamically adjusting the token bucket configuration for the subscriber based on the demand.

20 It is appreciated that the traffic regulator may take a variety of forms and includes a bucket mechanism. The bucket mechanism may also take a variety of forms and may include one or more token buckets, leaky buckets, or other bucket arrangements to provide more complex packet handling when required for a particular application. It is appreciated that the token bucket is an aspect of the
25 bucket mechanism behavior and may be achieved in any suitable way. That is, the bucket mechanism may have a number of aspects including a token bucket aspect to determine the conformance of incoming data. The bucket mechanism may also have additional aspects including, for example, a leaky bucket aspect to shape outgoing

data. It is appreciated that the various aspects of the bucket mechanism may be achieved through any suitable arrangements for processing data.

Handling packets in accordance with the token bucket configuration may occur in a variety of ways. In one aspect, the present invention comprehends
5 handling a packet in a normal fashion when sufficient tokens are present in the token bucket, and otherwise handling a packet in a special fashion. The special fashion used when sufficient tokens are not present in the bucket may involve dropping the packet or assigning a classification to the packet. In another aspect, the present invention comprehends handling a packet that arrives at the regulator by assigning
10 a classification to the packet that is based on the number of tokens present in the bucket when the packet arrives.

Implementations of the present invention that assign classifications to packets during packet handling (also sometimes called remarking) may treat the remarked packets in a number of different ways. The remarked packets may simply
15 be further processed by the bucket mechanism or sent after remarking with the remarking affecting downstream processing. Alternatively, in an implementation where the bucket mechanism includes multiple policy support, the remarked packets may be reprocessed by the bucket mechanism (because they are now part of a different class) so that the packets are handled according to the appropriate policy.

20 Further, in implementations of the present invention that assign classifications to packets during packet handling, the assigning of the classification may take place in accordance with any suitable classification assignment policy. For example, the assigning of the classification may take place in accordance with a predetermined relationship between the number of tokens present in the token bucket and the appropriate classification. For example, the assigning of the classification
25 may take place in accordance with a probability mass function that determines the probability mass for each classification based on a number of tokens present in the token bucket. When a probability mass function is used, a “drop” classification may be defined to designate packets that are to be dropped. In this way, the

probability mass function may provide random early (before the token bucket empties) packet dropping.

It is appreciated that measuring the demand placed on the packet communication network by the subscriber may occur in a variety of ways. In one aspect, the present invention comprehends measuring the demand by monitoring the number of tokens present in the token bucket. Preferably, measuring the demand further comprises determining a burst demand based on observations made while monitoring the number of tokens present in the token bucket over a period of time. Alternatively, the bucket mechanism may include a second bucket arrangement associated with the subscriber, and measuring the demand includes monitoring the second bucket arrangement. Further in the alternative, the demand may be measured by measuring the demand placed on the packet communication network by the subscriber for traffic in a different direction than the direction traffic that is regulated by the subscriber's token bucket.

It is appreciated that dynamically adjusting the token bucket configuration for the subscriber may occur in a variety of ways. In one aspect, the present invention comprehends dynamically adjusting the token bucket configuration by modifying the bucket depth. In an implementation that assigns classifications to packets during packet handling, dynamically adjusting may include modifying the policy to which the assigning of the classification conforms. The policy to which the assigning of the classification conforms may be, for example, based on a predetermined relationship between number of tokens present in the token bucket and appropriate classification, or based on a probability mass function that determines the probability mass for each classification (optionally including a "drop" classification) based on the number of tokens present in the token bucket.

The advantages associated with embodiments of the present invention are numerous. For example, the subscriber's token bucket configuration is dynamically adjusted based on demand to deliver higher performance for certain types of subscriber demand behavior. Adjustment may, for example, increase or decrease the effective burst ability (via, for example, changes to bucket depth and/or

classification assignment policy) to temporarily receive data at a rate exceeding the token bucket fill rate. Adjustment may, for example, ensure that bandwidth is equitably distributed among the subscribers in the face of network over-subscription, and in the face of vastly different network usage characteristics of average and heavy-usage subscribers.

Further, in carrying out the present invention, a method of traffic regulation in a packet communication network is provided. The network includes a traffic regulator for regulating traffic at the packet level. The traffic regulator includes a bucket mechanism. The bucket mechanism includes first and second token buckets associated with a subscriber. The first token bucket is configured to receive new tokens at a first fill rate and configured with a first bucket depth. The second token bucket is configured to receive new tokens at a second fill rate and has a second bucket depth. The method comprises handling packets that arrive at the regulator in accordance with the first and second token bucket configurations. The first token bucket uses tokens to regulate the packet flow in terms of packet rate. The second token bucket uses tokens to regulate the packet flow in terms of data rate. A particular packet is subjected to handling in accordance with both the first token bucket and the second token bucket.

Further, in carrying out the present invention, a method of traffic regulation in a packet communication network is provided. In this method, a token bucket is associated with a subscriber. Packets arriving at the regulator are handled in accordance with the token bucket configuration. The token bucket uses tokens to regulate the packet flow by removing tokens from the token bucket when handling packets. The amount of tokens to be removed is based on the amount of the flow and is further based on a classification of the flow. That is, the identical amount of flow may result in different amounts of tokens being removed for different classifications of traffic.

The amount of flow serving as the partial basis for the amount of tokens to be removed may be flow in terms of data or flow in terms of packets. The classification of the flow preferably dictates a token multiplier. That is, the amount

of the flow (in terms of data or in terms of packets depending on the implementation) determines a base number of tokens that is multiplied by the multiplier determined from the flow classification to determine the total number of tokens to be removed.

5 The present invention also comprehends packet-level devices programmed to carry out methods and techniques of the present invention. The present invention is particularly useful in a TCP/IP network or any other network with a connection-based protocol running over the packet-level traffic. In this way, dropped packets result in the connection-based protocol reducing the effective data
10 rate. In this way, the token bucket aspect of the bucket mechanism determines conformance of the incoming data and takes appropriate actions in the presence of non-conforming data.

 The above objects and other objects, features, and advantages of the present invention are readily apparent from the following detailed description of the
15 preferred embodiment when taken in connection with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

 FIGURE 1 is a cable network including a traffic regulator programmed to dynamically adjust the token bucket configuration for a subscriber based on demand in accordance with the present invention;

20 FIGURE 2 illustrates the traffic regulator in more detail;

 FIGURE 3 illustrates a first technique for packet handling;

 FIGURE 4 illustrates a second technique for packet handling;

 FIGURE 5 is a graph depicting classification assignment based on tokens present;

FIGURE 6 is a graph depicting classification assignment based on tokens present where the packets had an original class and are being remarked;

FIGURE 7 is a chart depicting classification assignment based on tokens present using a probability mass function;

5 FIGURE 8 illustrates a first technique for measuring the subscriber's network demand;

FIGURE 9 illustrates a second technique for measuring the subscriber's network demand;

10 FIGURE 10 illustrates a third technique for measuring the subscriber's network demand;

FIGURE 11 illustrates a dynamic adjustment of the token bucket configuration and is best understood with a comparison to Figure 5;

FIGURE 12 illustrates a dynamic adjustment of the token bucket configuration and is best understood with a comparison to Figure 6;

15 FIGURE 13 illustrates a dynamic adjustment of the token bucket configuration and is best understood with a comparison to Figure 7;

FIGURE 14 illustrates a traffic regulator in an alternative embodiment of the present invention that subjects packet flow to packet rate regulation and data rate regulation;

20 FIGURE 15 illustrates a traffic regulator in an alternative embodiment of the present invention that uses a token bucket based on amount and class of flow; and

FIGURE 16 illustrates using a class based multiplier to determine the amount of tokens to be removed from the token bucket during packet handling with the traffic regulator of Figure 15.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

5 With reference to Figure 1, a cable network is generally indicated at 10. Cable network 10 includes regional cable modem termination system (CMTS) 12. CMTS 12 includes a traffic regulator 14. CMTS 12 connects through hybrid fiber coax (HFC) network 16 to provide services to subscribers. The services available from CMTS 12 for subscriber 18 include Internet service as indicated at 10 20, telephony service as indicated at 22 and other content service as indicated at 24. Traffic regulator 14 attempts to manage network traffic at the packet level to avoid congestion and is illustrated in more detail in Figure 2..

 With reference to Figure 2, traffic regulator 14 regulates traffic at the packet level and includes bucket mechanism 30. Bucket mechanism 30 includes a 15 token bucket 32 associated with subscriber 18. Token bucket 32 is configured to receive new tokens at a fill rate and is configured with a bucket depth. It is appreciated that traffic regulator 14 may take a variety of forms and that bucket mechanism 30 may also take a variety of forms. Bucket mechanism 30 may include one or more token buckets, leaky buckets, or other bucket arrangements 34 to 20 provide more complex packet handling when required for a particular application. The configuration of token bucket 32 is dynamically adjusted based on network demand from the subscriber.

 Handling packets in accordance with the token bucket configuration of token bucket 32 may occur in a variety of ways. Figure 3 illustrates a first 25 technique for packet handling. Block 40 indicates packet arrival at token bucket 32. Decision block 42 checks to see if there are enough tokens present in token bucket 32 to send the packet normally. If sufficient tokens are present, the packet is sent normally at block 46. In the event that there are not enough tokens present in the bucket 32, the packet is either dropped or remarked as indicated at block 44.

Figure 4 illustrates a second technique for packet handling. Block 50 illustrates the arrival of the packet. Block 52 assigns a classification to the packet that is based on the number of tokens present in token bucket 32. At block 54, the packet is handled in accordance with the assigned classification.

5 Figure 5 is a graph depicting class assignment based on tokens present. In Figure 5, the x-axis represents the amount of tokens present in token bucket 32. The y-axis represents the assigned classification. As shown, the assigned classification is a higher class when more tokens are present in token bucket 32 and a lower class when less tokens are present in token bucket 32.

10 Figure 6 is a graph depicting class assignment based on tokens present where the packets arrived with an original class and are being remarked. The x-axis represents the amount of tokens present in token bucket 32. The y-axis represents the assigned classification. As shown, generally, the assigned classification is a higher class when more tokens are present in token bucket 32 and
15 a lower class when less tokens are present in token bucket 32. The relationship between number of tokens present in token bucket 32 and the assigned classification is different for each original class. As shown, packets originally marked as maximum class are always remarked at the highest assignable class. Packets originally marked as high class are remarked as high, medium, or low class
20 depending on the amount of tokens present in token bucket 32. Packets originally marked as medium class are remarked as medium or low class depending on the amount of tokens present in token bucket 32.

 When classifications are assigned to packets during packet handling, bucket mechanism 30 may treat the remarked packets in a number of different ways.
25 The remarked packets may simply be sent after remarking. Alternatively, in an implementation where bucket mechanism 30 includes multiple policy support, the remarked packets may be reprocessed by 30 bucket mechanism so that the packets are handled according to the appropriate policy.

Further, when classifications are assigned to packets during packet handling, the assigning of the classification may take place in accordance with any suitable classification assignment policy. Figures 5 and 6 illustrate the assigning of the classification in accordance with a predetermined relationship between the number of tokens present in the token bucket and appropriate classification. Figure 7 illustrates the assigning of the classification in accordance with a probability mass function that determines the probability mass for each classification based on a number of tokens present in the token bucket.

Figure 7 is a chart depicting classification assignment based on tokens present using a probability mass function. In Figure 7, which uses a left-handed coordinate system, the x-axis represents the amount of tokens present in token bucket 32. The y-axis represents the assigned classification. The z-axis represents probability mass. As shown, when many tokens are present, the high class has by far the greatest probability mass. As the number of tokens present decreases, the high class probability mass decreases, and the medium class probability mass increases. As the number of tokens present decreases even more, the high class probability mass continues to decrease, the medium class probability mass peaks and then decreases, and the low class probability mass increases rapidly.

A probability mass for a "drop" classification is also depicted on the chart in Figure 7. The "drop" classification designates packets that are to be dropped. The probability mass function provides random early (before the token bucket empties) packet dropping due to the "drop" class probability mass.

It is appreciated that the classification assignment using a probability mass function may also be used in applications where packets arrive with an original class and are remarked. In such a situation, each original class should have a corresponding probability mass function for assigning a class for remarking.

It is also appreciated that the traffic class for a particular packet may be indicated using any known technique for traffic classification. For example, known techniques involving DiffServe CodePoint (DSCP), Class of Service (CoS),

Quality of Service (QoS), and Type of Service (ToS) may be appropriate depending on the situation.

It is appreciated that measuring the demand placed on the network by the subscriber may occur in a variety of ways. Figure 8 illustrates a first technique for measuring the subscriber's demand. At block 60, the number of tokens present in token bucket 32 is monitored. At block 62, a burst demand is determined based on observations made while monitoring at block 60. Figure 9 illustrates a second technique for measuring the subscriber's demand. In this second technique, bucket mechanism 30 includes a second bucket arrangement associated with the subscriber (indicated by optional other bucket arrangements 34). At block 64, measuring the demand includes monitoring the second bucket arrangement. Figure 10 illustrates a third technique for measuring the subscriber's demand. In this third technique, demand is measured by measuring the demand placed on the network by the subscriber 18 for traffic in a different direction than the direction of traffic that is regulated by the subscriber's token bucket 32. The demand measurements may also be affected by time of day, overall network load, or other factors in addition to being affected by the activities of the subscriber.

It is appreciated that dynamically adjusting the token bucket configuration for the subscriber's token bucket 32 may occur in a variety of ways. The depth of token bucket 32 may be modified. In an implementation that assigns classifications to packets during packet handling, the policy to which the assigning of the classification conforms may be modified. Further, the token bucket configuration adjustment in accordance with the present invention comprehends adjustments that modify bucket depth and class assignments.

Figure 11 illustrates a dynamic adjustment of the token bucket configuration of Figure 5. As shown, going from the Figure 5 configuration to the Figure 11 configuration involves a change in the shape of the assignment classification function. Figures 5 and 11 do not specifically indicate amounts of tokens. In one form of the dynamic adjustment, going from the Figure 5 configuration to the Figure 11 configuration involves increasing the size of the token

bucket depth without changing the levels where assigned class changes. Accordingly, more overall burst capability is provided in Figure 11 and that additional part of the burst capability is at the highest class.

5 Figure 12 illustrates a dynamic adjustment of the token bucket configuration of Figure 6. As shown, going from the Figure 6 configuration to the Figure 12 configuration involves a change in the shape of the assignment classification function for each original class. Figures 6 and 12 do not specifically indicate amounts of tokens. In one form of the dynamic adjustment, going from the Figure 6 configuration to the Figure 12 configuration involves increasing the size
10 of the token bucket depth without changing the levels where assigned class changes (this is true for each original class). Accordingly, more overall burst capability is provided in Figure 12 (for each original class) and that additional part of the burst capability is at the highest class (for each original class).

Figure 13 illustrates a dynamic adjustment of the token bucket
15 configuration of Figure 7. As shown, going from the Figure 7 configuration to the Figure 13 configuration involves a change in the assignment classification probability mass function. Figures 7 and 13 do not specifically indicate amounts of tokens. In one form of the dynamic adjustment, going from the Figure 7 configuration to the Figure 13 configuration involves increasing the size of the token
20 bucket depth without changing the overall feel of the probability mass function for the previously existing domain. The added domain due to the increase in bucket depth has corresponding function values that provide additional overall burst capability particularly at the highest class.

In applications where packets arrive with an original class and are
25 remarked with a probability mass function, and each original class has a corresponding probability mass function for assigning a class for remarking, each probability mass function may be adjusted.

Figure 14 illustrates a traffic regulator 80 in an alternative embodiment of the present invention. Traffic regulator 80 includes bucket

mechanism 82. Bucket mechanism 82 includes first and second token buckets 84 and 86, respectively, and optional other bucket arrangements 88 to provide more complex packet handling when required for a particular application. First token bucket 84 is configured to receive new tokens at a first fill rate and has a first bucket depth. Second token bucket 86 is configured to receive new tokens at a second fill rate and has a second bucket depth. First token bucket 84 uses tokens to regulate the packet flow in terms of packet rate. Second token bucket 86 uses tokens to regulate the packet flow in terms of data rate. A particular packet is subjected to handling in accordance with both the first token bucket and the second token bucket. That is, when a packet arrives at bucket mechanism 82 for handling, in order to be sent normally, the packet must confirm to the requirements of token bucket 84 and token bucket 86. Put another way, the subscriber is simultaneously subjected to sustain rate and burst size limits in terms of packets and subjected to sustain rate and burst size limits in terms of data (bits or bytes). As further illustrated in Figure 14, token buckets 84 and 86 are subjected to dynamic adjustment based on network demand.

Figure 15 illustrates a traffic regulator 90 in another alternative embodiment of the present invention. Traffic regulator 90 includes bucket mechanism 92. Bucket mechanism 92 includes token bucket 94 and optional other bucket arrangements 96 to provide more complex packet handling when required for a particular application. Token bucket 94 uses tokens to regulate the packet flow by removing tokens from token bucket 94 when handling packets. The amount of tokens to be removed is based on the amount of the flow and is further based on a classification of the flow. That is, the identical amount of flow may result in different amounts of tokens being removed for different classifications of traffic. Token bucket 94 is subjected to dynamic adjustment based on network demand.

With reference to Figure 16, a preferred technique for determining the tokens to be removed from token bucket 94 to send a packet normally is illustrated. Token bucket 94 may regulate in terms of packets or in terms of data. Block 100 represents the determination of the base amount of tokens for the flow being processed. The classification of the flow in the preferred technique dictates

a token multiplier. Block 102 represents determining the multiplier. That is, the amount of the flow (in terms of data or in terms of packets depending on the implementation) determines a base number of tokens that is multiplied by the multiplier determined from the flow classification to determine (at block 104) the
5 total number of tokens to be removed.

There are many advantages associated with embodiments of the present invention. A subscriber's token bucket configuration is dynamically adjusted to modify aspects of the subscriber's network connection performance. These dynamic adjustments are made based on network demand measurements. In
10 this way, it is possible to deliver higher performance for certain types of subscriber behavior. The adjustments to the configuration may increase or decrease the subscriber's ability to receive bursts of data. Changes to bucket depth change the burst size while changes to class assignment policy change the way the data is handled further downstream. The token bucket configuration could be adjusted in
15 real-time to decrease a subscriber's burst capability when the subscriber engages in behavior that is generally sustain rate limited to prevent a continuous pattern of large bursts spaced apart by backoff periods. The token bucket could again be adjusted in real-time to increase the subscriber's burst capability when the subscriber ceases to engage in the sustain rate limited behavior. In this way, a subscriber is granted
20 very generous burst capability when the subscriber is not pushing the sustain rate limit for long periods of time. Accordingly, a subscriber engaging in basic low bandwidth web browsing or email would experience large bursts and perceive high performance. On the other hand, a subscriber engaging in demanding high bandwidth activities would have much more restricted burst capability which would
25 reduce network congestion.

While embodiments of the invention have been illustrated and described, it is not intended that these embodiments illustrate and describe all possible forms of the invention. Rather, the words used in the specification are words of description rather than limitation, and it is understood that various changes
30 may be made without departing from the spirit and scope of the invention.